

Apparatus and Method for Sending Point-to-point Protocol over Ethernet

This application is a continuation-in-part of application Serial Number 09/798,432, filed March 2, 2001, and which is incorporated herein in its entirety by reference.

Field of the Technology

This invention relates to computer networks and specifically to a method of transmitting and receiving information using point-to-point protocol ("PPP") over an ethernet network. Although the use of this invention is not limited to the internet, the internet provides the primary environment for practicing the invention.

Background of the Invention

The internet is not a single network, but comprises a large number of interconnected networks. When information is to be transmitted across the internet, the device originating the information, which may be a computer, will initially construct packets in which the data being transmitted is contained as a "payload." "Headers" and "trailers" conforming to the transmission protocols being used will be prepended and appended to the data to provide routers with sufficient information to forward the packets from network to network, in a series of "hops," until the packet arrives at its intended destination. As used in this specification, "packet" shall refer, generically, to a sequence of bytes representing a unit of data being transmitted pursuant to one or more transmission protocols. "Bytes" shall refer to an octet of binary digits. Once the packet arrives at its destination, headers and footers are stripped away, and the data is made available to the appropriate process running on the recipient computer.

Within the design of IP ("internet protocol"), every physical network has a maximum packet size, designated "maximum transmission unit," or MTU, and the MTU may be different for different networks. MTU is determined as a function of network design, including network bandwidth, maximal diameter, and desired imposed jitter. Since an IP packet in transit will frequently traverse more than a single network, it may encounter MTUs of different sizes. Since a packet cannot be transmitted over a network whose MTU is smaller than the packet size, one possible solution has been for a sending device to use a path MTU discovery algorithm to determine the smallest MTU that will be encountered during transit to the destination, and to establish a maximum packet size based upon that information. However, that solution has encountered a number of documented difficulties (RFC 2923, "TCP Problems with Path MTU Discovery"), and does not always present an acceptable solution for the problem.

Each network segment is defined by a router, and a packet passing through a router when transiting from one network to another will have its headers and trailers analyzed, stripped, modified, or added to by the router, depending upon the protocol being used by the next network segment. In order to route packets efficiently, routers maintain information about the networks connected to them, including the MTU. When a router encounters a packet that is larger than the MTU for the next network segment in the path to the packet's destination, the packet will not be accepted by the network segment, and may be lost, with a resulting communication failure between the sending and receiving devices. For this reason, it is important that packets be properly sized to be accepted by the networks they will be transiting.

1 Because each packet of information is discretely routed from source to
2 destination, packets may follow different paths, depending upon network conditions.
3 While most networks comprising the internet are high speed networks, using protocols
4 such as ATM and the like, conditions occasionally arise in which other, slower
5 transmission protocols and media are used. Under some circumstances, passage across a
6 network may involve a packet's being transmitted across an ethernet network using point-
7 to-point protocol ("PPP"). Such protocols may be found in dial-up networks, ISDN, and,
8 more recently, DSL networks, and are frequently used to connect individual devices to an
9 internet service provider. When this combination of protocols is used, it is not
10 uncommon for difficulties to arise that culminate in the loss of transmitted data.

11 Data to be transmitted to a remote device will normally be generated by a process
12 running on a computer. The data will be sent to a TCP buffer in the RAM of the
13 computer where it will be formatted and encapsulated within a TCP header and an IP
14 header which provide addressing information for the packet and for the process on the
15 recipient machine. Thereafter, additional headers will be added, depending upon the
16 network protocols being used on the network to which the computer is connected. For
17 ethernet networks, the last header to be added will be an ethernet header, which is added
18 by the ethernet driver that is attached to the physical transmission medium. When the
19 packet is received at the destination, a reverse process is employed to decapsulate the
20 packet and provide data to the appropriate process running on the destination computer.
21 The processes of encapsulation and decapsulation, and associated functions of receiving,
22 comparing, setting option and header values, transmitting, and the like, are carried out by
23 programs and drivers running on the sending device.

Ethernet is a low-level network protocol, and is the primary protocol found in local area networks (LANs). Ethernet frames transport data carried in higher level protocols across ethernet networks. Ethernet drivers accept information formatted by upper level protocols such as IP, TCP (transmission control protocol), ARP (address resolution protocol), and ICMP (internet control message protocol), and “encapsulate” it for delivery across the ethernet network.

Ethernet is a multiple access network in which many devices may be attached to the same physical transmission medium. Because each device on an ethernet network must be able to be uniquely distinguished from the others, each is identified by a globally unique physical address, sometimes referred to as a “medium access control”, or “MAC” address. When information is to be delivered across an ethernet network, the sending device adds an eight byte preamble and an ethernet header at the beginning of the packet. The ethernet header is 14 bytes, and comprises the destination device’s MAC address, the sending device’s MAC address, and the ethernet type. A 4-byte trailer comprising a check frame sequence is appended to the packet. The packet is then transmitted to the network, and the device that recognizes its own address in the destination address field receives the frame.

Ethernet frames may be of varying length. However, the maximum permissible length of an ethernet frame which, by convention, does not include the preamble, but which does include the header (which holds the source and destination addresses, and the ethernet type), and the trailing Frame Check Sequence, is 1518 bytes.

Information formatted in higher level protocols, such as IP, TCP, or PPP, is contained in a data field, or “payload,” that is located between the ethernet frame’s

1 The basic TCP header is 20 bytes in length, and relies upon the IP header within
2 which it is encapsulated to provide source and destination device addresses. The TCP
3 header includes source and destination ports, and other information needed to place
4 packets in sequence, to control packet fragmentation, to acknowledge receipt of a packet,
5 to verify the integrity of information, to signal various conditions, and to carry out other
6 functions. The TCP header may also contain options which will control the handling of
7 following TCP packets in the session. One of those options is a maximum segment size
8 ("MSS") value which occupies 4 bytes of the TCP options field (2 bytes identify the
9 option as MSS and two bytes represent the number of bytes for the maximum segment
10 size). When set, this number limits the number of bytes in the TCP payload that the
11 sending device is prepared to receive throughout the session.

12 The header of a TCP packet for "opening" a socket for communications will set a
13 flag bit to signal a SYN (synchronize) condition, and will include other information that
14 is used in the session associated with the socket being opened. The MSS value can be set
15 only in the initial SYN packet. Other options, such as the Window Scale option and the
16 SACK ("selective acknowledgment) are also available only in an initial SYN packet.
17 Once the TCP session has been opened, and throughout the session until the session is
18 closed (by setting a bit in the FIN flag) the TCP parameters for communicating with the
19 socket will remain as they were established when the session was opened, and the TCP
20 header will remain at a constant length of 20- bytes throughout the session.

21 The point-to-point protocol ("PPP") is a set of interdependent protocols designed
22 to work together to support the concurrent operation of multiple higher-layer protocols
23 over a PPP serial link. PPP is an IETF (Internet Engineering Task Force) Standard

specified in RFC-1661. PPP provides a standard for transporting such higher-level protocols between two peer devices by encapsulating higher-level data along with negotiation mechanisms for configuring the link. The PPP header may include configuration options, one of which is a "maximum-receive-unit" (MRU). This option may be sent to inform the peer (receiving device) that the implementation can receive larger packets, or to request that the peer send smaller packets. The default MRU is 1500 bytes.

PPP is probably best known for use in telephone or ISDN dial-up links, or DSL connections between individual computers and internet service providers ("ISPs") who provide a connection to the internet. Data formatted for IP is encapsulated within a PPP packet for delivery from the individual computer to the ISP. At the ISP, the encapsulation will be stripped away, and the IP packet will be delivered to the internet for further transmission to its destination.

Because PPP was developed as a protocol to connect two "peer" devices, it lends itself to methods of access control, billing functionality, and type of service demands. These features and controls, although desirable under particular circumstances, are specific to "two-party" networks, and are not available in traditional ethernet networks. These desirable features of PPP have led to recent efforts to develop a method for transmitting PPP over ethernet networks. These efforts are described in RFC-2516 which, although not an internet standard, proposes a method for transmitting PPP over Ethernet ("PPPoE") by encapsulating PPP packets within ethernet packets to provide many of the benefits associated with each of the protocols.

1 The PPPoE header for an ethernet frame is 6 bytes long. The payload of a PPPoE
2 packet includes a PPP packet, whose header is 2 bytes in length, and any other packets
3 that may be encapsulated within the PPP packet. Optional "tags" attached to the PPPoE
4 packet are carried in the payload section, and may further reduce the maximum PPP
5 payload size. In order to accommodate the PPP packet within the ethernet frame, RFC
6 2516 provides that the MRU option must not be negotiated to be larger than 1492 bytes.
7 This options is relevant, however, only when the PPP packet will be received by the
8 device that will generate a responding transmission. However, when the packet that is
9 encapsulated within the PPP packet is destined for a device that lies beyond the network
10 segment that is using PPP, the PPP and PPPoE headers will be stripped from the packet
11 before it reaches its destination, and the packet will then be routed to its final destination
12 without the MRU information. When this happens, the receiving machine will not be
13 aware that the packet it sends in response will be transiting a network segment using PPP
14 protocol on its trip back to the sending device, and it will default to sending a packet
15 whose size is limited to the maximum size for an ethernet payload, or 1500 bytes.

16 When this responding packet reaches the router immediately preceding the PPPoE
17 segment, the addition of the PPP (2 byte) and PPPoE (6 byte) headers may increase the
18 size of the ethernet payload to more than 1500 bytes, if the payload's original size had
19 been larger than 1492 bytes. When that happens, the packet will be larger than the MTU
20 for that network, will not be able to transit the network segment, and will be lost.

21 The method and apparatus of the present invention uses the initializing TCP
22 header to carry information to the receiving machine to limit the size of TCP packets
23 transmitted from the receiving device to the sending device. This ensures that packets

Description of the Preferred Embodiment

Figure 1 depicts a hypothetical network having three network segments. A first computer 2 is located at one end, while a second computer 4 is located at the other end. The three network segments are connected by routers 6 and 8. Depictions of a single packet of information are shown at each network segment. When the packet is sent from the first computer 2, it is traversing a network segment that uses point-to-point protocol over ethernet. This may typically be a DSL connection from a home or office to an internet service provider. The packet 10 has a TCP packet that is encapsulated within an IP packet which, in turn, is encapsulated within a PPP packet. The PPP packet is encapsulated within a PPPoE packet, which itself is encapsulated within an ethernet packet. In accordance with the present invention, as the packet left the sending computer 2, the MSS option field value was set at "1452" bytes. In addition, the MRU option of the PPP packet would have been set at 1492. If the PPP were being used on a serial network having only two devices, the receiving device would be able to use the MRU to send responding packets of the requested size. In Figure 1, however, the packet 10 will be received at router 6, and will be routed to router 8 on an ethernet segment that does not

optional 4 byte MSS is 24-bytes in length. In this packet the SYN flag 130 would be set, indicating that a session is being initiated and a socket is being opened for interprocess communications. The TCP packet has a payload 120 whose maximum size is determined by the MSS value in the TCP options field 110. The TCP payload 120 carries process-specific information from a socket in the sending device to a corresponding socket in the receiving device. A 4-byte trailing frame check sequence (FCS) 140 is appended to the ethernet packet.

The MSS is a 16 bit number that theoretically may be as large as 65,535. However, because the maximum size for an ethernet payload (not including the ethernet header or trailer) is 1500 bytes, it is clear that any packet in which the size of the ethernet packet, including both the 14 byte header and the 4 byte file check sequence, exceeds 1518 bytes cannot be transmitted over an ethernet medium.

In order to limit ethernet packet length when using PPP, the preferred embodiment of this invention initializes a TCP session by substituting the number "1452" (0x05ac in hexadecimal) into the MSS field when the SYN flag 130 is set in the TCP header. This is shown in Figure 2 at 110. The value of 1452 is determined by subtracting from the maximum payload value for an ethernet frame (1500 bytes) the number of bytes in the headers of the encapsulated packets. These are, the PPPoE header (6 bytes), the PPP header(2 bytes), the IP header (20 bytes) and the TCP header (20 bytes).

Within a TCP header, the MSS field is one of the options that must be included in a TCP packet to open a socket for a session. Any such TCP socket opening packet may be identified by the SYN flag 130 in the header, which is set for socket opening frames and otherwise is clear. None of the optional fields, including the MSS, the window scale

option or the SACK options, will be needed for later transmissions once the session has started.

Figure 3 shows an ethernet packet in which PPP is encapsulated, and the TCP header does not include an options field. Because this packet does not open a session, the SYN flag 130 in the TCP header is clear. For non-initializing TCP packets, the TCP payload will always be preceded by the basic 20 byte TCP header.

The method of this invention can be implemented through software or firmware in any PPPoE session. Implementation may take the form of checking the MSS value for any TCP SYN packet and replacing any MSS value with "1452" if the original MSS value is larger than 1452; or the method could simply write the number "1452" into the MSS field for each TCP SYN packet, without first analyzing the existing value.

Although the preferred embodiment substitutes the value "1452" into the MSS option for initializing TCP packets, those of skill in the art will appreciate that any other number that is less than 1452 may be substituted into the MSS field, and will ensure that the receiving device will send responding packets that are more than 8 bytes smaller than the maximum size for an ethernet packet. Other network factors may indicate the use of a smaller packet size, although a smaller packet size may require more packets to be transmitted to convey the same data, resulting in a decrease in communications efficiency. It will be understood that the description herein relates to the preferred embodiment of the invention, and that the scope of the invention will encompass a range of MSS values, and is limited only by the following claims.